arXiv:1704.02306v1 [math.NT] 7 Apr 2017

COMPUTING ZETA FUNCTIONS OF GENERIC PROJECTIVE HYPERSURFACES IN LARGER CHARACTERISTIC

JAN TUITMAN

Abstract. We give improvements of the deformation method for computing the zeta function of a generic projective hypersurface in characteristic p that either reduce the dependence on p of the time complexity to $\tilde{O}(p^{1/2})$ or that of the space complexity to $\tilde{O}(\log(p))$ while remaining polynomial in the other input parameters.

Contents

1. Introduction 1
2. The deformation method 3
3. Our new algorithms 5
3.1. Diagonal fibre 6
3.2. Differential equation 7
References 10

1. Introduction

Let $\mathbf{F}_q$ be a finite field of characteristic p and cardinality $q = p^a$. Suppose that X is a projective hypersurface of degree d in projective n-space $\mathbf{P}^n_{\mathbf{F}_q}$. For all $i \in \mathbf{N}$ we can then count the number of points of X with values in $\mathbf{F}_{q^i}$. Note that X is defined by some homogeneous polynomial of degree d in $\mathbf{F}_q[x_0, \ldots, x_n]$ and $X(\mathbf{F}_{q^i})$ is just the number of zeros of this polynomial in projective n-space over $\mathbf{F}_{q^i}$. These numbers of points can be used to define a generating series that is called the zeta function of X:

$$Z(X,T) = \exp\left(\sum_{i=0}^{\infty} |X(\mathbf{F}_{q^i})| \frac{T^i}{i}\right).$$

It is well known that $Z(X,T)$ is the quotient of two polynomials with integer coefficients, so can be given by a finite amount of data. An obvious question is therefore whether it can be computed effectively. This is clearly the case, since there are well known bounds for the degrees of the numerator and denominator of $Z(X,T)$, so that one is reduced to computing a finite number of the $X(\mathbf{F}_{q^i})$, which can be determined by naive counting. A more interesting problem is whether, and if so how, $Z(X,T)$ can be computed efficiently.

Lauder and Wan [13] proposed an algorithm that computes $Z(X,T)$ in time $(pad^n)^{O(n)}$ without any additional assumptions on X. Abbott, Kedlaya and Roe [1] proposed a different algorithm with roughly the same complexity (but a smaller constant) wich assumes X to be generic (at least smooth). The reason for this assumption is that in contrast to Lauder and Wan, the algorithm of Abbott, Kedlaya and Roe uses p-adic cohomology, which is better behaved in the smooth case.

Note that the input size of the problem, i.e. the bit-size of the defining polynomial of X, is about $\log(p)ad^n$. Therefore, the algorithms from the previous paragraph are not polynomial time for multiple reasons: p^n is not polynomial in $\log(p)$, a^n is not polynomial in a and d^{n^2} is not polynomial in d^n. Only for fixed p and n (but varying a and d) are these algorithms polynomial time.

Lauder [11, 12] then introduced the deformation method, which computes $Z(X,T)$ in time $(pad^n)^{O(1)}$, for a generic hypersurface X. Note that this is a polynomial time algorithm for

fixed p. The main idea is to compute the p-adic cohomology of X by deforming it to a diagonal hypersurface. The author and Pancratz [14] improved the deformation method further and showed that it can be made to run in

$$\text{time: } \tilde{O}\left(pa^3d^{n(\omega+4)}e^{n(\omega+1)}\right), \qquad \text{space: } \tilde{O}\left(pa^3d^{5n}e^{3n}\right),$$

where ω denotes an exponent for matrix multiplication, e is the basis of the natural logarithm and we use the $\tilde{O}(-)$ notation that ignores logarithmic factors, i.e. $\tilde{O}(f)$ denotes the class of functions that lie in $O(f\log^k(f))$ for some $k \in \mathbf{N}$.

All algorithms mentioned so far run in time and space at least linear in p. Note that this is typically the case for all p-adic point counting algorithms e.g. [10, 4, 15, 16]. Recently, Harvey has improved this situation, first for hyperelliptic curves [7, 8] and then for general schemes [9]. One algorithm from [9] runs in

$$\text{time: } O\left(p\log(p)^{1+\epsilon}2^{6n^2+13n}a^{3n+4+\epsilon}n^{3n+4+\epsilon}(d+1)^{3n^2+6n+\epsilon}\right),$$
$$\text{space: } O\left(\log(p)2^{4n^2+9n}a^{2n+3}n^{2n+2}d^{2n^2+4n}\right),$$

and another one in

$$\text{time: } O\left(p^{1/2}\log(p)^{2+\epsilon}2^{8n^2+16n}a^{4n+4+\epsilon}n^{4n+4+\epsilon}(d+1)^{4n^2+7n+\epsilon}\right),$$
$$\text{space: } O\left(p^{1/2}\log(p)2^{4n^2+9n}a^{2n+3}n^{2n+2}d^{2n^2+4n}\right).$$

Note that both of these algorithms do not require the hypersurface X to be generic, since the use of p-adic cohomology is avoided.

We see that Harvey's algorithms are better than the deformation method in terms of the dependence of the time and/or space complexity on p, but quite a lot worse in terms of the dependence on everything else (i.e. a, d, n). In particular, Harvey's algorithms are not polynomial time for fixed p anymore, since the complexity is polynomial in d^{n^2} instead of d^n and in a^n instead of a. The goal of this paper is to modify the deformation method so that it runs either in time $\tilde{O}(p^{1/2})$ or space $\tilde{O}(\log(p))$ in p, while remaining polynomial in a, d^n. More precisely, we will prove the following theorem.

Theorem 1.1. Let X be a generic hypersurface of degree d in projective space $\mathbf{P}^n_{\mathbf{F}_q}$ over a finite field $\mathbf{F}_q$ of characteristic p not dividing d and cardinality $q = p^a$. Then the zeta function $Z(X,T)$ may be computed in

$$\text{time: } \tilde{O}\left(pa^3d^{n(\omega+4)}e^{n(\omega+1)}\right), \qquad \text{space: } \tilde{O}\left(\log(p)a^3d^{5n}e^{3n}\right),$$

or alternatively in

$$\text{time: } \tilde{O}\left(p^{1/2}a^3d^{n(2\omega+3)}e^{n(\omega+1)}\right), \qquad \text{space: } \tilde{O}\left(p^{1/2}a^3d^{n(2\omega+1)}e^{3n}\right).$$

Proof. This is a simplified version of Theorem 3.7 below. □

What it means for a hypersurface to be generic is explained in Remark 2.11 below. Comparing our new algorithms to Harvey's we see that our exponents of a and d^n are constant instead of $O(n)$ and smaller for any value of n. The dependence on p is more or less the same (probably exactly the same, with some more work). The only remaining advantages of Harvey's algorithms over ours are that they can be applied more generally and are conceptually simpler.

Another algorithm from [9] starts from a hypersurface X of degree d in projective n-space $\mathbf{P}^n_{\mathbf{Z}}$ over the integers defined by a homogeneous polynomial $f \in \mathbf{Z}[x_0, \ldots, x_n]$ and computes $Z(X \otimes \mathbf{F}_p, T)$ for all primes $p < N$ in

$$\text{time: } O\left(N\log^2(N)\log^{1+\epsilon}(N\|f\|)2^{8n^2+16n}n^{4n+6+\epsilon}(d+1)^{4n^2+7n+\epsilon}\right),$$
$$\text{space: } O\left(N\log(N)\log(ndN\|f\|)2^{4n^2+11n}n^{2n+4}(d+1)^{2n^2+5n}\right),$$

where $\|f\|$ denotes a bound on the coefficients of f. Since there are about $N/\log(N)$ primes $p < N$, the time per prime is actually polynomial in $\log(N)$, so that this is an average polynomial

time algorithm. We expect that using the ideas in this paper and the formalism of accumulating remainder trees [6, 8], one should be able to obtain an algorithm that is average polynomial time in p and polynomial time in a and d^n at the same time. However, because of some technical complications, for now this remains work in progress.

The author was supported by FWO Vlaanderen. We thank David Harvey for helpful discussions.

2. The deformation method

In this section we briefly recall how the deformation method works. For more details the reader should consult [14]. We will keep the terminology and notation from that paper here as much as possible.

Let $\mathbf{F}_q$ be a finite field of characteristic p and cardinality $q = p^a$. Suppose that X_1 is a projective hypersurface of degree d in projective n-space $\mathbf{P}^n_{\mathbf{F}_q}$ and let $U_1 = \mathbf{P}^n_{\mathbf{F}_q} \setminus X_1$ denote its complement. Moreover, let $\mathbf{Q}_q$ be the unique unramified extension of $\mathbf{Q}_p$ with residue field $\mathbf{F}_q$, let $\mathbf{Z}_q$ be its ring of integers and $\sigma \in \mathrm{Gal}(\mathbf{Q}_q/\mathbf{Q}_p)$ the unique lift of the p-th power map on $\mathbf{F}_q$. We extend σ to the standard p-th power Frobenius lift on $\mathbf{P}^1_{\mathbf{Q}_q}$, i.e. the one that sends the standard coordinate t to t^p. Finally, let $\mathrm{ord}_p(-)$ denote the p-adic valuation on $\mathbf{Q}_q$.

For an algebraic variety Y over $\mathbf{F}_q$, let $H^i_{\mathrm{rig}}(Y)$ denote the rigid cohomology spaces of Y. These are finite dimensional vector spaces over $\mathbf{Q}_q$ that are contravariantly functorial in Y, and are equipped with a σ-semilinear action of the p-th power Frobenius map on Y that we denote by F_p. For the construction and basic properties of these spaces we refer to [2].

To compute the zeta function $Z(X_1, T)$ it is sufficient to compute the cohomology space $H^n_{\mathrm{rig}}(U_1)$ with the action of F_p by the following theorem.

Theorem 2.1. We have

$$Z(X_1, T) = \frac{\chi(T)^{(-1)^n}}{(1-T)(1-qT)\cdots(1-q^{n-1}T)},$$

where $\chi(T) = \det\big(1 - T(p^{-1}\,\mathrm{F}_p)^a | H^n_{\mathrm{rig}}(U_1)\big) \in \mathbf{Z}[T]$ denotes the reverse characteristic polynomial of the action of $(p^{-1}\,\mathrm{F}_p)^a$ on $H^n_{\mathrm{rig}}(U_1)$. Moreover, the polynomial $\chi(T)$ has degree

$$\frac{1}{d}\big((d-1)^{n+1} + (-1)^{n+1}(d-1)\big).$$

Proof. [14, Theorem 2.13]. □

The main idea of the deformation method is to take advantage of the way the rigid cohomology spaces and the action of F_p vary in a family.

Definition 2.2. Let $P_1 \in \mathbf{Z}_q[x_0, \ldots, x_n]$ be any homogeneous polynomial of degree d that defines a hypersurface $\mathcal{X}_1 \subset \mathbf{P}^n_{\mathbf{Z}_q}$ such that $\mathcal{X}_1 \otimes \mathbf{F}_q \cong X$ and let $P_0 \in \mathbf{Z}_q[x_0, \ldots, x_n]$ be a diagonal polynomial $P_0 = a_0 x_0^d + \ldots + a_n x_n^d$ with $a_1, \ldots, a_n \in \mathbf{Z}_p^\times$. Define $P \in \mathbf{Z}_q[t][x_0, \ldots, x_n]$ to be

$$P = (1-t)P_0 + tP_1.$$

The polynomial P defines a family $\mathcal{X}/\mathcal{S}$ of smooth projective hypersurfaces over some Zariski open subset $\mathcal{S}$ of $\mathbf{P}^1_{\mathbf{Z}_q}$ such that the fibre $\mathcal{X}_0$ at $t = 0$ is the diagonal hypersurface defined by P_0 and the fibre $\mathcal{X}_1$ at $t = 1$ is isomorphic to $\mathcal{X}_1$. Let $\mathcal{U}/\mathcal{S}$ denote the complement of $\mathcal{X}/\mathcal{S}$. We denote the generic fibres of $\mathcal{X}, \mathcal{S}, \mathcal{U}$ by $\mathfrak{X} = \mathcal{X} \otimes \mathbf{Q}_q$, $\mathfrak{S} = \mathcal{S} \otimes \mathbf{Q}_q$, $\mathfrak{U} = \mathcal{U} \otimes \mathbf{Q}_q$ and the special fibres by $X = \mathcal{X} \otimes \mathbf{F}_q$, $S = \mathcal{S} \otimes \mathbf{F}_q$, $U = \mathcal{U} \otimes \mathbf{F}_q$, respectively.

Theorem 2.3. The vector bundle $H^n_{\mathrm{dR}}(\mathfrak{U}/\mathfrak{S})$ with its Gauss–Manin connection ∇ admits a Frobenius structure F with the following property: for any $\tau \in S(\mathbf{F}_q)$ with Teichmüller lift $\hat{\tau} \in \mathcal{S}(\mathbf{Z}_q)$, we have

$$(H^n_{\mathrm{rig}}(U_\tau), \mathrm{F}_p) \cong (H^n_{\mathrm{dR}}(\mathfrak{U}/\mathfrak{S}), F)_{\hat{\tau}}$$

as $\mathbf{Q}_q$-vector spaces with a σ-semilinear endomorphism. We will therefore denote this Frobenius structure on $H^n_{\mathrm{dR}}(\mathfrak{U}/\mathfrak{S})$ by F_p as well.

Proof. [14, Theorem 2.10]. □

We can explicitly write down a basis for $H^n_{\mathrm{dR}}(\mathfrak{U}/\mathfrak{S})$, at least generically.

Definition 2.4. For $k \in \mathbf{N}$, we define the following sets of monomials:

$$B_k = \{x^u \colon u \in \mathbf{Z}^{n+1}_{\geq 0}, |u| = kd - (n+1) \text{ and } u_i < d-1 \text{ for all } i\},$$

where $x^u = x_0^{u_0} \dots x_n^{u_n}$ and $|u| = u_0 + \dots u_n$. If x^u is contained in some B_k, then we let $k(u)$ denote this k. Let Ω denote the n-form on $\mathfrak{U}/\mathfrak{S}$ defined by

$$\Omega = \sum_{i=0}^{n} (-1)^i x_i dx_0 \wedge \dots \wedge \widehat{dx_i} \wedge \dots \wedge dx_n,$$

where $\widehat{dx_i}$ means that dx_i is left out. We then define $\mathcal{B}_k = \{Q\Omega/P^k \colon Q \in B_k\}$, $B = B_1 \cup \dots \cup B_n$, $\mathcal{B} = \mathcal{B}_1 \cup \dots \mathcal{B}_n$ and $|B| = b$.

Theorem 2.5. The set $\mathcal{B}$ (restricted to the fibre at $t = 0$) forms a basis for $H^n_{\mathrm{rig}}(U_0)$. Moreover, $\mathcal{B}$ also forms a basis for $H^n_{\mathrm{dR}}(\mathfrak{U} \otimes \mathbf{Q}_q(t))$, so that the same conclusion holds for almost all other fibres.

Proof. [14, Theorem 3.9]. □

Remark 2.6. Note that $b = \frac{1}{d}\big((d-1)^{n+1} + (-1)^{n+1}(d-1)\big)$ by Theorem 2.1.

Let $M \in M_{b\times b}(\mathbf{Q}_q(t))$ denote the matrix of the Gauss–Manin connection ∇ with respect to the basis $\mathcal{B} = [\omega_1, \dots, \omega_b]$, i.e.

$$\nabla(\omega_j) = \sum_{i=1}^{b} M_{ij}\omega_i \otimes dt$$

and let $r \in \mathbf{Z}_q[t]$ with $\mathrm{ord}_p(r) = 0$ denote a denominator for M, i.e. such that we can write $M = G/r$ with $G \in M_{b\times b}(\mathbf{Q}_q[t])$. Moreover, let Φ denote the matrix of $p^{-1}\,\mathrm{F}_p$ with respect to the basis $\mathcal{B}$, i.e.

$$p^{-1}\,\mathrm{F}_p(\omega_j) = \sum_{i=1}^{b} \Phi_{ij}\omega_i.$$

Note that Φ has entries in the ring

$$\mathbf{Q}_q\,\langle t, 1/r\rangle^\dagger = \left\{ \sum_{i,j=0}^{\infty} a_{i,j}\frac{t^i}{r^j} \;:\; a_{i,j} \in \mathbf{Q}_q,\ \exists c > 0 \text{ s.t. } \lim_{i+j\to\infty} \big(\mathrm{ord}_p(a_{i,j}) - c(i+j)\big) \geq 0 \right\},$$

of overconvergent functions. The fact that F defines a Frobenius structure on the vector bundle $H^n_{\mathrm{dR}}(\mathfrak{U}/\mathfrak{S})$ with its Gauss–Manin connection ∇ translates into the following differential equation for the matrix Φ.

Theorem 2.7. The matrix $\Phi \in M_{b\times b}(\mathbf{Q}_q\langle t, 1/r\rangle^\dagger)$ satisfies the differential equation

$$\left(\frac{d}{dt} + \frac{G}{r}\right)\Phi = pt^{p-1}\Phi\sigma(M), \qquad \Phi(0) = \Phi_0,$$

where Φ_0 is the matrix of $p^{-1}\,\mathrm{F}_p$ on $H^n_{\mathrm{rig}}(U_0)$ with respect to the basis $\mathcal{B}$.

Proof. [14, Theorem 2.17]. □

The matrix Φ_0 can be computed using an explicit formula from [14] that we will now recall.

Definition 2.8. For $l \in \mathbf{Q}$ and $r \in \mathbf{Z}_{\geq 0}$, let the rising factorial $\prod_{j=0}^{r-1}(l+j)$ be denoted by $(l)_r$.

Definition 2.9. Let $u, v \in \mathbf{Z}^{n+1}$ be such that we have $x^u, x^v \in B$ and $p(u_i+1) \equiv v_i + 1 \bmod d$ for all i. We define

$$\alpha_{u,v} = \prod_{i=0}^{n} a_i^{(p(u_i+1)-(v_i+1))/d} \left(\sum_{m,r} \left(\frac{u_i+1}{d}\right)_r \sum_{j=0}^{r} \frac{\left(pa_i^{p-1}\right)^{r-j}}{(m-pj)!j!} \right),$$

where the sum in the i-th factor of the product is over all integers $m, r \geq 0$ that satisfy the equation $p(u_i+1) - (v_i+1) = d(m-pr)$.

Theorem 2.10. Let ω_i denote an element of $\mathcal{B}$ corresponding to a tuple $u \in \mathbf{Z}^{n+1}$ and let ω_j denote the unique element of $\mathcal{B}$ corresponding to a tuple $v \in \mathbf{Z}^{n+1}$ such that $p(u_l + 1) \equiv v_l + 1 \bmod d$ for all l. Then we have

$$p^{-1} \operatorname{F}_p(\omega_i) = (-1)^{k(v)} \frac{(k(v)-1)!}{(k(u)-1)!} p^{n-k(u)} \alpha_{u,v}^{-1} \cdot \omega_j \tag{2.1}$$

as elements of $H^n_{\mathrm{rig}}(U_0)$.

Proof. [14, Theorem 4.3]. □

The deformation method consists of the following steps:

Step 1. Compute the matrix $M \in M_{b\times b}(\mathbf{Q}_q(t))$ of the Gauss–Manin connection ∇.

Step 2. Compute the matrix $\Phi_0 \in M_{b\times b}(\mathbf{Q}_p)$ of the action of the map $p^{-1}\operatorname{F}_p$ on $H^n_{\mathrm{rig}}(U_0)$ using Theorem 2.10.

Step 3. Solve the differential equation from Theorem 2.7 for $\Phi \in M_{b\times b}(\mathbf{Q}_q\langle t, 1/r\rangle^\dagger)$ and evaluate at $t = 1$ to obtain the matrix $\Phi_1 \in M_{b\times b}(\mathbf{Q}_q)$ of the action of the map $p^{-1}\operatorname{F}_p$ on $H^n_{\mathrm{rig}}(U_1)$.

Step 4. Use Theorem 2.1 to deduce the zeta function $Z(X_1, T)$.

Remark 2.11. We have already mentioned that the deformation algorithm only applies to generic hypersurfaces. We will now explain what the exact conditions are [14].

First, the hypersurfaces $\mathcal{X}_0$ and $\mathcal{X}_1$ have to be smooth over $\mathbf{Z}_q$. For $\mathcal{X}_0$ this is guaranteed if p does not divide d and for $\mathcal{X}_1$ if the hypersurface X_1 is smooth.

Second, we need that the matrix M of the Gauss–Manin connection does not have a pole in the open p-adic unit disk around $t = 1$. Note that such a pole can only be an apparent singularity, caused by the generic basis $\mathcal{B}$ not being a basis at the pole. By [14, Proposition 3.13], it is sufficient that $\overline{R}(1) \neq 0$, where $\overline{R} \in \mathbf{F}_q[t]$ is the reduction modulo p of the polynomial $R \in \mathbf{Z}_q[t]$ defined in [14, Definition 3.12]. When X_1 varies, $\overline{R}(1)$ is a polynomial in the coefficients of its defining polynomial. Since $\overline{R}(1) \neq 0$ for diagonal hypersurfaces X_1 by the proof of [14, Theorem 3.6], we have that $\overline{R}(1) \neq 0$ for generic X_1.

Therefore, even for fixed $a_0, \ldots, a_n \in \mathbf{Z}_p^\times$ defining P_0, as long as p does not divide d the deformation method can be applied to a generic hypersurface X_1.

Theorem 2.12. In the deformation method as presented in [14], the four steps above have the following complexities [14, §7]:

Step 1.	time $\tilde{O}(\log(p) a^2 (d^{n(\omega+2)} e^{n(\omega+1)} + d^{5n} e^{3n}))$	space $\tilde{O}(\log(p) a^2 d^{4n} e^{3n})$
Step 2.	time $\tilde{O}(p a^3 d^{4n})$	space $\tilde{O}(\log(p) a d^{2n})$
Step 3.	time $\tilde{O}(p a^3 d^{n(\omega+4)} e^{2n})$	space $\tilde{O}(p a^3 d^{5n} e^n)$
Step 4.	time $\tilde{O}(\log^2(p) a^2 d^{n(\omega+1)})$	space $\tilde{O}(\log(p) a^2 d^{3n})$

Proof. [14, §7], where we have taken $a = a'$ and $d_t = 1$, because of our choice of P. □

Clearly the only steps not polynomial in $\log(p)$ are the second and the third, both of which are quasilinear in p. So to improve the dependence of the complexity on p, we only have to consider these steps. All the p-adic and t-adic precisions required to obtain provably correct results remain the same as in [14].

3. Our new algorithms

Our strategy to reduce the dependence of the complexity on p to $\tilde{O}(p^{1/2})$ will be the same as in [7]: reduce the steps in our algorithm that are quasilinear in p to computing certain matrix products and then apply an algorithm of Bostan, Gaudry and Schost [3] (based on [5]) to compute these products faster than using the naive method. The precise result that we need is as follows.

Theorem 3.1. Let R be a commutative ring and M (respectively MM) a function $\mathbf{N} \to \mathbf{N}$ such that polynomials of degree less than d (respectively matrices of size $m \times m$) can be multiplied in $\mathsf{M}(d)$ (respectively $\mathsf{MM}(m)$) ring operations (i.e. $+, -, \times$) in R. Moreover, let $A(x)$ be an $m \times m$ matrix over the polynomial ring $R[x]$ of degree at most 1. Suppose that the invertibility

conditions of [3, Theorem 14] are satisfied. Then for any positive integer N, the matrix product $A(1)A(2)\cdots A(N)$ can be computed in

$$O\left(\mathsf{MM}(m)\sqrt{N} + m^2\mathsf{M}(\sqrt{N})\right)$$

ring operations in R, storing only $O(m^2\sqrt{N})$ elements of R.

Proof. This is a special case of [7, Theorem 9], which is based on [3, Theorem 15]. □

Remark 3.2. We do not explain the invertibility conditions in this theorem in detail, since for us R will always be one of the p-adic fields $\mathbf{Q}_p$ or $\mathbf{Q}_q$, for which they are trivially satisfied.

3.1. **Diagonal fibre.**

We start by analyzing Step 2 of the deformation method, i.e. the computation of the matrix Φ_0. In [14, §4] it is explained that we have to compute the

$$\alpha_{u,v} = \prod_{i=0}^{n} a_i^{(p(u_i+1)-(v_i+1))/d}\left(\sum_{m,r}\left(\frac{u_i+1}{d}\right)_r \sum_{j=0}^{r}\frac{\left(pa_i^{p-1}\right)^{r-j}}{(m-pj)!j!}\right), \tag{3.1}$$

to p-adic precision $N'_{\Phi_0} \in \tilde{O}(ad^n)$ and that for this we only have to consider terms in the outer sum with $r \le \mathcal{R}, m \le \mathcal{M}$ and $p(u_i+1)-(v_i+1) = d(m-pr)$, where $\mathcal{R} \in \tilde{O}(ad^n)$ and $\mathcal{M} \in O(p\mathcal{R})$. Note that all p-adic precisions in the algorithm are $\tilde{O}(ad^n)$ by [14, §7], so that a single multiplication in $\mathbf{Q}_p$ takes time and space $\tilde{O}(\log(p)ad^n)$.

From the formula above we see that we have to compute $k!$ (to finite p-adic precision) for all $i = 0, \ldots, n$ and all $k \le \mathcal{M}$ such that

$$k \equiv \frac{p(u_i+1)-(v_i+1)}{d} \bmod p.$$

Computing these factorials naively as in [14] clearly takes time at least linear in p, so we will have to proceed differently.

Let us fix a basis vector $u \in B$ and let $v \in B$ be the unique basis vector such that we have $p(u_i+1)-(v_i+1) \equiv 0 \bmod d$ for all $0 \le i \le n$. For a single value of $0 \le i \le n$ we can compute and store all $k!$ with $k \le \mathcal{M}$ and

$$k \equiv \frac{p(u_i+1)-(v_i+1)}{d} \bmod p$$

in time and space $\tilde{O}(p^{1/2}a^2d^{2n})$. Indeed, we need to compute $k!$ for $\mathcal{R} \in \tilde{O}(ad^n)$ different values of k and to go from one value of k to the next we have to multiply by the product

$$(k+1)\ldots(k+p),$$

which can be computed in time and space $\tilde{O}(p^{1/2}ad^n)$ using the algorithm from Theorem 3.1 with $m = 1$, $R = \mathbf{Q}_p$, $A(x) = (x+k)$ and $N = p$. Doing this computation for all $0 \le i \le n$ instead of a single value of i does not change the time and space requirements, since the factor n is absorbed by the $\tilde{O}$ symbol. This takes care of all the $(m-pj)!$ we have to compute in (3.1).

The elements

$$a_i^{(p(u_i+1)-(v_i+1))/d}, \left(\frac{u_i+1}{d}\right)_r, (pa_i^{p-1})^{r-j} \text{ and } j!$$

that we need in (3.1) (note that u is still fixed), can be computed and stored naively in time and space

$$\tilde{O}(n\mathcal{R}\log^2(p)ad^n) \subset \tilde{O}(\log^2(p)a^2d^{2n}).$$

Moreover, $\alpha_{u,v}$ can be computed from these elements and the $k!$ in time

$$\tilde{O}(n\mathcal{R}^2\log(p)ad^n) \subset \tilde{O}(\log(p)a^3d^{3n})$$

and additional space $\tilde{O}(\log(p)ad^n)$. Putting everything together, we find that a single $\alpha_{u,v}$ can be computed in time $\tilde{O}(p^{1/2}a^2d^{2n} + a^3d^{3n})$ and space $\tilde{O}(p^{1/2}a^2d^{2n})$. Since there are $b \in O(d^n)$ different basis vectors u, all of the $\alpha_{u,v}$ can be computed in time $\tilde{O}(p^{1/2}a^2d^{3n} + a^3d^{4n})$ and space $\tilde{O}(p^{1/2}a^2d^{2n})$. The same is then clearly true for the matrix Φ_0 by Theorem 2.10.

Summarising, we have proved the following theorem:

Theorem 3.3. Step 2 of the deformation method, i.e. the computation of the matrix Φ_0 to the required p-adic precision can be carried out in

$$\text{time:} \quad \tilde{O}(p^{1/2}a^2d^{3n} + a^3d^{4n}), \qquad \text{space:} \quad \tilde{O}(p^{1/2}a^2d^{2n}).$$

3.2. **Differential equation.**

We now move on to the third step of the deformation method, i.e. the computation of the matrix Φ_1 by solving the differential equation from Theorem 2.7 for Φ and evaluating this matrix at $t = 1$. Let us start by recalling what we did in [14], to understand why that is quasilinear in p and how it can be improved.

Recall from [14, §5] that if $C \in M_{b\times b}(\mathbf{Q}_q[[t]])$ denotes a fundamental matrix of horizontal sections of the Gauss–Manin connection ∇ with respect to the basis $\mathcal{B}$, i.e. such that

$$\left(\frac{d}{dt} + \frac{G}{r}\right) C = 0, \qquad C(0) = I,$$

then $\Phi = C\Phi_0\sigma(C^{-1})$ is a solution to the differential equation from Theorem 2.7. Let us write $G = \sum_{i=0}^{\deg(G)} G_i t^i$ and $r = \sum_{i=0}^{\deg(r)} r_i t^i$. Note that $\deg(G), \deg(r) \in \tilde{O}((de)^n)$ by [14, §7]. A power series solution $C = \sum_{i=0}^{\infty} C_i t^i$ can be obtained by solving the following matrix recurrence [14, (5.5)]:

$$C_0 = I,$$

$$C_{i+1} = \frac{-1}{r_0(i+1)} \left(\sum_{j=i-\deg(G)}^{i} G_{i-j}C_j + \sum_{j=i-\deg(r)+1}^{i} r_{i-j+1}(jC_j) \right),$$

where we take $C_j = 0$ for $j < 0$.

The inverse matrix C^{-1} satisfies the dual differential equation [14, Remark 5.7]:

$$\left(\frac{d}{dt} - \frac{G^t}{r}\right) (C^{-1})^t = 0, \qquad C^{-1}(0) = I,$$

which can be solved by a similar recurrence.

To obtain Φ as an element of $M_{b\times b}(\mathbf{Q}_q[[t]])$ to t-adic precision K, we need to determine C_i for $i < K$ and $(C^{-1})_i$ for $i < \lceil K/p \rceil$. Then we compute the product

$$\Phi = C\Phi_0\sigma(C^{-1}),$$

where Φ_0 has already been obtained in the second step. It is shown in [14, §7] that to recover Φ_1 to the required p-adic precision, we have to take

$$K \in \tilde{O}(pad^{2n}e^n).$$

Moreover, before we can evaluate Φ at $t = 1$, we have to multiply it by some polynomial $s \in \mathbf{Z}_q[t]$ [14, Theorem 6.6] with $\deg(s) \in \tilde{O}(pad^{2n}e^n)$ [14, §7] and truncate to t-adic precision K again, because of some technical convergence issues. Finally, we evaluate $(s/s(1))\Phi$ at $t = 1$ to obtain the matrix Φ_1. Exact bounds for all p-adic and t-adic precisions can be found in [14, §5, §6]. All p-adic precisions in the algorithm are $\tilde{O}(ad^n)$, so elements of $\mathbf{Q}_q$ can be multiplied in time and space $\tilde{O}(\log(p)a^2d^n)$.

The computation from the previous paragraph takes time at least linear in p: the number of terms of C and s is K so computing these objects or multiplying by them already has complexity at least linear in p. Note that we do not have this problem for the matrix C^{-1} since that has $\lceil K/p \rceil$ terms which is $\tilde{O}(ad^{2n}e^n)$ by Theorem 3.4 below. Recall that σ sends t to t^p, so that the number of nonzero terms in $\sigma(C^{-1})$ to t-adic precision K is still $\tilde{O}(ad^{2n}e^n)$. Finally, the evaluation of $(s/s(1))\Phi$ at $t = 1$ again has complexity at least linear in p, because there are K terms. Hence we will have to proceed differently.

We may assume $s = r^\theta$ with $\theta/p \in \tilde{O}(ad^n)$, by Theorem 3.4 below. We have to compute:

$$(s/s(1))\Phi = (r/r(1))^\theta C\Phi_0\sigma(C^{-1})$$

as an element of $M_{b\times b}(\mathbf{Q}_q[[t]])$ to t-adic precision K and then evaluate at $t = 1$. We know that multiplying $(r/r(1))^\theta$ and C has complexity at least linear in p. However, we can avoid this

multiplication by solving a slightly different differential equation instead. Let $D \in M_{b\times b}(\mathbf{Q}_q[[t]])$ denote the matrix $D = (r/r(1))^\theta C$. Then D satisfies the differential equation

$$\left(\frac{d}{dt} + \frac{G - \theta\frac{dr}{dx}I}{r}\right) D = 0, \qquad\qquad D(0) = \left(\frac{r(0)}{r(1)}\right)^\theta I.$$

So we can avoid the multiplication by $(r/r(1))^\theta$ simply by replacing the matrix G by the matrix

$$H = G - \theta\frac{dr}{dx}I$$

and solving the recurrence

$$D_0 = \left(\frac{r(0)}{r(1)}\right)^\theta I,$$
$$D_{i+1} = \frac{-1}{r_0(i+1)}\left(\sum_{j=i-\deg(H)}^{i} H_{i-j}D_j + \sum_{j=i-\deg(r)+1}^{i} r_{i-j+1}(jD_j)\right),$$

where we take $D_j = 0$ for $j < 0$.

The matrix Φ_1 is now given by the evaluation of $D\Phi_0\sigma(C^{-1})$ to t-adic precision K at $t = 1$, which we can also write as

$$\Phi_1 = \sum_{j=0}^{\lceil K/p\rceil - 1}\left(\sum_{i=0}^{K-1-pj} D_i\right)\cdot\Phi_0\cdot\sigma((C^{-1})_j).$$

The number of terms in the inner sum (say for j=0) is at least linear in p, so evaluating it will have complexity at least linear in p. However, since the D_i satisfy a recurrence, so do the sums $E_i = D_0 + \ldots + D_i$:

$$E_0 = \left(\frac{r(0)}{r(1)}\right)^\theta I,$$
$$E_{i+1} = E_i + \frac{-1}{r_0(i+1)}\left(\sum_{j=i-\deg(H)}^{i} H_{i-j}(E_j - E_{j-1}) + \sum_{j=i-\deg(r)+1}^{i} r_{i-j+1}\cdot j\cdot(E_j - E_{j-1})\right),$$

where we take $E_j = 0$ for $j < 0$.

The expression for the matrix Φ_1 becomes

$$\Phi_1 = \sum_{j=0}^{\lceil K/p\rceil - 1} E_{K-1-pj}\cdot\Phi_0\cdot\sigma((C^{-1})_j). \tag{3.2}$$

Note that the orders of the recurrences for $C_i, (C^{-1})_i, D_i, E_i$ are all $\tilde{O}((de)^n)$ since $\deg(G)$ and $\deg(r)$ are $\tilde{O}((de)^n)$. However, to apply Theorem 3.1, we need a recurrence of order 1. The order of the recurrence for E_i is given by

$$\kappa = \max\{\deg(H) + 1, \deg(r)\} + 1.$$

Let $\mathbb{E}_i \in M_{\kappa b\times b}(\mathbf{Q}_q)$ be the matrix with vertical blocks $E_i, E_{i-1}, \ldots, E_{i-\kappa+1}$. Then from the recurrence for E_i, we find a matrix $A(x) \in M_{b\times b}(\mathbf{Q}_q[x])$ of degree 1 such that

$$\mathbb{E}_{i+1} = (i+1)^{-1}A(i+1)\mathbb{E}_i$$

for all $i \geq 0$. Note that E_i can be easily read off from $\mathbb{E}_i$.

Recall that $\lceil K/p\rceil \in \tilde{O}(ad^{2n}e^n)$. Hence the $(C^{-1})_i$ with $i \leq \lceil K/p\rceil$ can be computed and stored in

$$\text{time: } \tilde{O}((K/p)(de)^n b^\omega a^2 d^n) \subset \tilde{O}(\log(p)a^3 d^{n(\omega+4)}e^{2n}),$$
$$\text{space: } \tilde{O}((K/p)b^2 a^2 d^n) \subset \tilde{O}(\log(p)a^3 d^{5n}e^n).$$

Applying σ to C^{-1} takes time

$$\tilde{O}((K/p)b^2(\log^2(p)a + \log(p)a^2 d^n)) \subset \tilde{O}(\log^2(p)a^3 d^{5n}e^n)$$

and no additional space. This is all the same as in [14, Proposition 7.7]. Now to compute Φ_1, we can proceed in two different ways.

The first way is to evaluate (3.2) by solving the recurrence for E_i naively. Note that we only have to store the last $\tilde{O}((de)^n)$ matrices E_i. Therefore, this takes

$$\text{time: } \tilde{O}(K(de)^n b^\omega a^2 d^n) \subset \tilde{O}(pa^3 d^{n(\omega+4)} e^{2n}),$$
$$\text{space: } \tilde{O}((de)^n b^2 \log(p) a^2 d^n) \subset \tilde{O}(\log(p) a^2 d^{4n} e^n).$$

The second way is to determine $\mathbb{E}_{i+p}$ from $\mathbb{E}_i$ repeatedly, using Theorem 3.1 to compute

$$A(i+p)A(i+p-1)\dots A(i+1) \qquad \text{and} \qquad (i+1)(i+2)\dots(i+p).$$

That the indices in the first product are decreasing instead of increasing can easily be circumvented by tranposing. Note that we will have to increase the p-adic precision by

$$\operatorname{ord}_p((i+1)(i+2)\dots(i+p)) \in O(\log_p(K)),$$

but this remains $\tilde{O}(ad^n)$. Therefore $\mathbb{E}_{i+p}$ can be determined from $\mathbb{E}_i$ in time and space

$$\tilde{O}(p^{1/2}(\kappa b)^\omega a^2 d^n) \subset \tilde{O}(p^{1/2} a^2 d^{n(2\omega+1)} e^{n\omega}).$$

We need to do this $\lceil K/p \rceil$ times and update Φ_1 each time. This takes time

$$\tilde{O}\left((K/p) p^{1/2} a^2 d^{n(2\omega+1)} e^{n\omega} + (K/p) b^\omega \log(p) a^2 d^n\right) \subset \tilde{O}\left(p^{1/2} a^3 d^{n(2\omega+3)} e^{n(\omega+1)}\right)$$

and no additional space.

In the discussion above we have used two results that are implicit in [14, §7], but cannot be found there explicitly. For completeness we prove these results in the next theorem.

Theorem 3.4. We can take:

(i) $s = r^\theta$ with $\theta/p \in \tilde{O}(ad^n)$ and $\deg(r) \in \tilde{O}((de)^n)$,

(ii) $K/p \in \tilde{O}(ad^{2n} e^n)$.

Proof. We can take r equal to the polynomial $R = \prod_{i=2}^{n+1} \det(\Delta_k)$ from [14, Definition 3.12] which has degree $\tilde{O}((de)^n)$. Let N_Φ denote the p-adic precision to which the matrix $\Phi \in M_{b\times b}(\mathbf{Q}_q\langle t, 1/r\rangle)$ has to be computed. From [14, §7], we know that $N_\Phi \in \tilde{O}(ad^n)$. By [14, Theorems 6.4 (with $z \neq \infty$)] we can take

$$s = \det(\Delta_2 \dots \Delta_n) \det(\Delta_{n+1})^{p(n+h(N_\Phi))-n},$$

where

$$h(N_\Phi) = \max\left\{i \in \mathbf{N} \colon i + (n-1) + \operatorname{ord}_p((n-1)! - n\lfloor \log_p(p(n+i)-n)\rfloor < N_\Phi\right\}.$$

We can take $\theta = p(n + h(N_\Phi))$, so that $\theta/p \in \tilde{O}(ad^n)$ (note that the logarithm in the definition of $h(N_\Phi)$ is to base p). By [14, Theorem 6.4 (with $z = \infty$)] it follows that

$$K \leq \deg(s) + 1 + \deg\Big(\prod_{k=2}^{n} \Delta_k^{-1}\Big) + (p(n + h(N_\Phi)) - n)\deg(\Delta_{n+1}^{-1}) + ph(N_\Phi),$$

where deg denotes minus the order at $z = \infty$, so that $K/p \in \tilde{O}(ad^{2n} e^n)$. □

Summarising, we have proved the following theorem.

Theorem 3.5. Step 3 of the deformation method, i.e. the computation of the matrix Φ_1 to the required p-adic precision, can be carried out in either

$$\text{time: } \tilde{O}\left(pa^3 d^{n(\omega+4)} e^{2n}\right), \qquad \text{space: } \tilde{O}\left(\log(p) a^3 d^{5n} e^n\right),$$

or alternatively

$$\text{time: } \tilde{O}\left(p^{1/2} a^3 d^{n(2\omega+3)} e^{n(\omega+1)}\right), \qquad \text{space: } \tilde{O}\left(p^{1/2} a^2 d^{n(2\omega+1)} e^{n\omega} + a^3 d^{5n} e^n\right).$$

Proof. This follows by adding up all complexities from the discussion above, leaving out terms that are dominated by other terms or powers of logarithms of other terms. The two different bounds correspond to the two different ways of evaluating Φ_1 explained above. □

Remark 3.6. Note that if $K/p \in O(K^{1/2-\epsilon})$ for some $0 < \epsilon < 1/2$, then by [7, Theorem 9] (which is slightly stronger than Theorem 3.1), all the way at the end of the $\tilde{O}(p^{1/2})$ algorithm we can save a factor K/p and compute Φ_1 to the required p-adic precision in

$$\text{time: } \tilde{O}\left(p^{1/2}a^2d^{n(2\omega+1)}e^{n\omega} + a^3d^{n(\omega+3)}e^n\right), \quad \text{space: } \tilde{O}\left(p^{1/2}a^2d^{n(2\omega+1)}e^{n\omega} + a^3d^{5n}e^n\right).$$

This improvement only applies when p is large enough relative to a and d^n. Since the exact condition on p is rather complicated, we will not use this improvement in the theorem below.

The total complexity of our new algorithms is as follows.

Theorem 3.7. Let X be a generic hypersurface of degree d in projective space $\mathbf{P}^n_{\mathbf{F}_q}$ over a finite field $\mathbf{F}_q$ of characteristic p not dividing d and cardinality $q = p^a$. Then the zeta function $Z(X,T)$ may be computed in

$$\begin{aligned}&\text{time: } \tilde{O}\left(pa^3d^{n(\omega+4)}e^{2n} + a^2(d^{n(\omega+2)}e^{n(\omega+1)} + d^{5n}e^{3n})\right),\\ &\text{space: } \tilde{O}\left(\log(p)\left(a^3d^{5n}e^n + a^2d^{4n}e^{3n}\right)\right),\end{aligned}$$

or alternatively in

$$\begin{aligned}&\text{time: } \tilde{O}\left(p^{1/2}a^3d^{n(2\omega+3)}e^{n(\omega+1)}\right),\\ &\text{space: } \tilde{O}\left(p^{1/2}a^2d^{n(2\omega+1)}e^{n\omega} + a^3d^{5n}e^n + a^2d^{4n}e^3n\right).\end{aligned}$$

Proof. This follows from Theorems 2.12, 3.3 and 3.5. □

Remark 3.8. Work on a Magma implementation of the algorithms in this paper is in progress and a first version should be out soon.

References

[1] T.G. Abbott, K.S. Kedlaya, and D. Roe. Bounding Picard numbers of surfaces using p-adic cohomology. In *Arithmetic, Geometry, and Coding Theory (AGCT-10)*, 2006.

[2] P. Berthelot. Géométrie rigide et cohomologie des variétés algébriques de caractéristique p. *Mém. Soc. Math. Fr. (N.S.)*, 23, 1986.

[3] Alin Bostan, Pierrick Gaudry, and Éric Schost. Linear recurrences with polynomial coefficients and application to integer factorization and Cartier-Manin operator. *SIAM J. Comput.*, 36(6):1777–1806, 2007.

[4] W. Castryck, J. Denef, and F. Vercauteren. Computing zeta functions of nondegenerate curves. *IMRP Int. Math. Res. Pap.*, pages Art. ID 72017, 57, 2006.

[5] D. V. Chudnovsky and G. V. Chudnovsky. Approximations and complex multiplication according to Ramanujan. In *Ramanujan revisited (Urbana-Champaign, Ill., 1987)*, pages 375–472. Academic Press, Boston, MA, 1988.

[6] Edgar Costa, Robert Gerbicz, and David Harvey. A search for Wilson primes. *Math. Comp.*, 83(290):3071–3091, 2014.

[7] David Harvey. Kedlaya's algorithm in larger characteristic. *Int. Math. Res. Not. IMRN*, (22):Art. ID rnm095, 29, 2007.

[8] David Harvey. Counting points on hyperelliptic curves in average polynomial time. *Ann. of Math. (2)*, 179(2):783–803, 2014.

[9] David Harvey. Computing zeta functions of arithmetic schemes. *Proc. Lond. Math. Soc. (3)*, 111(6):1379–1401, 2015.

[10] K.S. Kedlaya. Counting points on hyperelliptic curves using Monsky-Washnitzer cohomology. *Journal of the Ramanujan Mathematical Society*, 16, 2001.

[11] A.G.B. Lauder. Counting solutions to equations in many variables over finite fields. *Foundations of Computational Mathematics*, 4(3):221–267, 2004.

[12] A.G.B. Lauder. Deformation theory and the computation of zeta functions. *Proc. London Math. Soc*, 3:565–602, 2004.

[13] A.G.B. Lauder and D. Wan. Counting points on varieties over finite fields of small characteristic. In *Algorithmic number theory: lattices, number fields, curves and cryptography*, volume 44 of *Math. Sci. Res. Inst. Publ.*, pages 579–612. Cambridge Univ. Press, Cambridge, 2008.

[14] Sebastian Pancratz and Jan Tuitman. Improvements to the deformation method for counting points on smooth projective hypersurfaces. *Found. Comput. Math.*, 15(6):1413–1464, 2015.

[15] Jan Tuitman. Counting points on curves using a map to $\mathbf{P}^1$. *Math. Comp.*, 85(298):961–981, 2016.

[16] Jan Tuitman. Counting points on curves using a map to $\mathbf{P}^1$, II. *Finite Fields and Their Applications*, 45:301 – 322, 2017.

KU Leuven, Departement Wiskunde, Celestijnenlaan 200B, 3001 Leuven, Belgium
E-mail address: jan.tuitman@kuleuven.be